arXiv:1401.6628v1 [cs.DC] 26 Jan 2014

BigOP: Generating Comprehensive Big Data Workloads as a Benchmarking Framework

Yuqing Zhu, Jianfeng Zhan, Chuliang Weng♯, Raghunath Nambiar♢, Jinchao Zhang, Xingzhen Chen, and Lei Wang

State Key Laboratory of Computer Architecture (Institute of Computing Technology, Chinese Academy of Sciences), ♯Huawei, ♢Cisco
{zhuyuqing, zhanjianfeng, zhangjinchao, chenxingzhen, wanglei_2011}@ict.ac.cn,
♯chuliang.weng@huawei.com, ♢RNambiar@cisco.com

Abstract. Big Data is considered proprietary asset of companies, organizations, and even nations. Turning big data into real treasure requires the support of big data systems. A variety of commercial and open source products have been unleashed for big data storage and processing. While big data users are facing the choice of which system best suits their needs, big data system developers are facing the question of how to evaluate their systems with regard to general big data processing needs. System benchmarking is the classic way of meeting the above demands. However, existent big data benchmarks either fail to represent the variety of big data processing requirements, or target only one specific platform, e.g. Hadoop.

In this paper, with our industrial partners, we present Big***OP***, an end-to-end system benchmarking framework, featuring the abstraction of representative ***O***peration sets, workload ***P***atterns, and prescribed tests. BigOP is part of an open-source big data benchmarking project, ***BigDataBench***[1]. BigOP's abstraction model not only guides the development of BigDataBench, but also enables automatic generation of tests with comprehensive workloads.

We illustrate the feasibility of BigOP by implementing an automatic test generation tool and benchmarking against three widely used big data processing systems, i.e. Hadoop, Spark and MySQL Cluster. Three tests targeting three different application scenarios are prescribed. The tests involve relational data, text data and graph data, as well as all operations and workload patterns. We report results following test specifications.

1 Introduction

Companies, organizations and countries are taking big data as their important assets, as the era of big data has inevitably arrived. But drawing insights from big data and turning big data into real treasure demand an in-depth extraction of its values, which heavily relies upon and hence boosts the deployment of massive big data systems.

Big data owners are facing the problem of how to choose the right system for their big data processing requirements, while a variety of commercial and open

[1] BigDataBench is available at http://prof.ict.ac.cn/BigDataBench

source products, e.g., NoSQL databases [9], Hadoop MapReduce [5], Spark [28], Impala [3], Hive [7] and Redshift [1], have been unleashed for big data storage and processing. On the other hand, big data system developers are in need of application-perspective evaluation methods for their systems. Benchmarking is the classic way to direct the evaluation and the comparison of systems.

Though many well-established benchmarks exist, e.g. TPC series benchmarks [13] and HPL benchmarks [8], no widely accepted benchmark exists for big data systems. Some benchmarks targeting *big data systems* appear in recent years [22,23,18,21], but they are either for a specific platform or covering limited workload patterns.

Together with our industrial partners, we present in this paper an end-to-end system benchmarking framework BigOP, which enables automatic generation of tests with comprehensive workloads for big data systems. We build BigOP for our urgent need to benchmark big data systems. BigOP is part of a comprehensive big data benchmarking suite ***BigDataBench***[27], which is already used by our collaborators in testing architecture, network and energy efficiency of big data systems. The development of BigDataBench is guided by BigOP's abstraction model.

BigOP features an abstracted set of ***O***perations and ***P***atterns for big data processing. We work out the abstraction after considering the powerful representativeness of the five primitive relational operators [17] and the 13 computation patterns summarized in a report by a multidisciplinary group of well-known researchers [14]. The operations are extended from the five primitive relational operators, while the workload patterns are summarized based on general big data computation characteristics.

Figure 1 demonstrates an overview of BigOP. In BigOP, a benchmarking test is specified as a prescription for one application or a range of applications. A prescription includes a subset of operations and processing patterns, a data set, a workload generation method, and the metrics. The subset of operations and processing patterns are selected from BigOP's whole abstraction set. The data set can be obtained from real applications or generated through widely-obtainable tools. The workload generation method describes how operations are issued from clients, e.g, the number of client threads, the load, etc. The metrics can include the test duration, request latency metrics, and the throughput. With BigOP, a prescribed test can be implemented over different systems for comparison. System users can also prescribe a test targeting their specific applications.

BigOP leaves the choice of data set to users because the *variety* of big data makes a predefined data set for benchmarking irrelevant. Besides, data sets are usually related to the processing performance in big data scenarios, for example, highly isolated webpages vs. highly linked webpages for PageRank computations. The size of the chosen data set is required to be larger than the total memory size of the system under test (SUT) so that the volume characteristic of big data is covered. The velocity characteristic of big data can be represented in the workload generation specification. That is, BigOP design takes the three properties of Big Data into consideration.

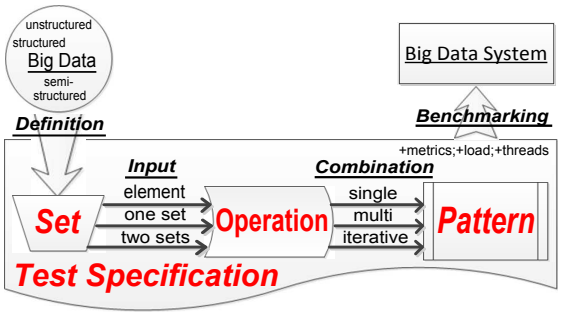

Fig. 1. Overview of BigOP.

After presenting the design of BigOP (**Section 2**), we give three test prescriptions targeting three different application scenarios as an example. We benchmark against three widely used big data processing systems, i.e. MySQL cluster, Hadoop[5]+HBase[6] and Spark [28], using BigOP (**Section 3**). The tests involve relational data, text data and graph data, as well as all operations and workload patterns. We discuss workload representativeness by comparing YCSB [18], TPC-DS [26] and BigBench [22] to BigOP. We summarize related work (**Section 4**) and conclude (**Section 5**) in the end.

2 BigOP Design

2.1 Overview

BigOP is an end-to-end system benchmarking framework with a big data processing operation and pattern (***OP***) abstraction. Comprehensive workloads can be specified and thus constructed based on the ***OP*** abstraction.

An adequate level of abstraction and the end-to-end execution model leaves space for various system implementations and optimizations. Benchmarks with these two properties enable comparisons among different kinds of systems servicing the same goals. The success of TPC benchmarks [13] demonstrates this fact [16]. Therefore, BigOP takes an end-to-end benchmarking model. Benchmarking workloads are applied by clients issuing requests through interfaces. Requests are sent through a network to the system under test (SUT). Metrics are measured at the client side.

The success of TPC benchmarks also highlights the importance of benchmarking systems with functions of abstraction and the functional workload model [16]. Functions of abstraction are basic units of computation occurring frequently in applications, while the functional workload model includes functions of abstraction, the representative application load, and the data set. Functions of abstraction is the core. To generalize functions of abstraction for BigOP, we abstract a set of operations and patterns common to big data processing.

Furthermore, big data embodies great variety. So do big data applications. Therefore, BigOP allows users to flexibly specify data sets and workloads in test prescriptions. The test prescription allows for application-specific features, as well as comparisons across systems.

Table 1. Basic Operations

Categories	Typical Operations
Element Operation	put, get, delete; transform; filter
Single-Set Operation	project, order by aggregation(min,max,sum,median,average)
Double-Set Operation	union, difference, cross product

2.2 Operation and Pattern Abstraction for Big Data Processing

Before going into the details of BigOP's operation and pattern abstraction, we first review some facts about big data processing.

The concept of *set* in relational algebra [17] is still effective in big data processing scenarios, e.g. the MapReduce [19] model, which plays an important role in big data processing. In the model, a large piece of data is transformed into a set of elements denoted by key-value pairs in the *map* stage. The *reduce* stage does further processing over the mapped set. We thus adopt the general concept of *set* in BigOP. Elements in a set can be uniquely identified.

Data accesses to memory and disk, as well as across network, must be considered in big data benchmarking. Memory size plays an important role as for system performance. As technology improves, the starting data size of TPC-H increases along with the obtainable amount of memory. Thus, benchmarks must consider the whole system, including the system composition. Besides, big data systems can consist of not only nodes, but also datacenters, due to the huge volume of big data. Thus, communication must also be considered in benchmarks. Furthermore, the huge volume of big data and the resulting processing complexity demand distribution and parallelization of computation tasks. Otherwise, the time required for big data processing would be intolerably long. Hence, BigOP requires the data set size to exceed the total memory size of the SUT, but a single element in the data set must fit into a single node's memory to be processed; or, the element can at least be read serially and transformed into a set of smaller elements for processing.

Operation Abstraction Considering the above facts, we first abstract big data processing operations into three categories, i.e., *element operation*, *single-set operation* and *double-set operation*. Element operation can be computed based on an individual element, which might require only local memory access. Single-set operations are computed based on elements of one set. Double-set operations require input from two sets of data. We did not include multi-set operations, which can be composed by combining multiple double-set operations. The more sets are involved in a processing task, the more demands for data accesses involving memory, disk and network communication are there. Operations from the three categories can be combined and permutated to meet complex processing requirements. Table 1 illustrates the three categories of operations.

BigOP adopts the five primitive operators from *relational algebra* [17], which also takes a *set*-based perspective. They are filter(select), project, cross product, set union, and set difference. The five primitive operators are fundamental in the sense that omitting any of them causes a loss of expressive power. Many other set operations can be defined in terms of these five. The *filter* operation is

Table 2. Workload Patterns

Patterns	Example Workloads
Single-Operation Processing	any abstracted operation
Multi-Operation Processing	operation combinations, SQL queries
Iterative Processing	graph traversal, finite state machines

an *element* operation because it can operate a set element on given conditions with only element-local information. The *project* operation is in the *single-set* operation category, requiring the set information. The *union*, *difference*, and *cross product* fall in the *double-set* operation category.

The basic data access operations of *put*, *get* and *delete* are in the element operation category. We also add a *transform* operation to this category because it is common to turn a big element into a set of elements or another element, as demonstrated by the Map usages of MapReduce. *transform* is user-defined. Most big data are unstructured data, therefore *transform* is important to define data-specific computations.

We also include the commonly used *order by* and *aggregation* operations in the single-set operation category. *order by* is equal to *sort*, a fundamental database operation noted by Jim Gray [14]. Quite a few benchmarks have been built based on *sort* [11]. *Aggregation* is included because it is widely recognized important operation to turn sets into numerals.

Pattern Abstraction The abstracted operations can be combined into more complex processing tasks following some patterns. We summarize three patterns as demonstrated with examples in Table 2. The three patterns are *single-operation*, *multi-operation* and *iterative* processing patterns. The single-operation pattern contains only a single operation in a processing task, while the multi-operation and iterative patterns can have multiple operations in a task. The inclusion of multiple operations allows the big data system to make a whole optimization plan for all operations in the task. The difference between multi-operation processing and iterative processing patterns is whether the exact number of operations to be executed is known beforehand. Iterative processing patterns only provide stopping conditions (which can be specified as user-defined functions), thus the exact number of operations can only be figured out in running time.

Different from SQL queries, the processing patterns in Table 2 can result in more than one set. Furthermore, the element definition of a data set relies on the *transform* operation, instead of *schema*. For example, a text document can be *transformed* into a set with *word* elements or with *sentence* elements. While element operation can be processed locally, global optimization techniques can be employed for single-set and double-set operations, as well as for multi-operation and iterative processing patterns.

The choice of the operations in Table 1 is in no way complete, but it is representative enough to represent a broad range of processing workloads when used with the patterns of Table 2. Besides, we think the efforts in benchmarking should be incremental and evolving. That is, more basic operations can be added to Table 1 in the future, as well as more patterns to Table 2.

Table 3. Test Prescriptions without Workload Generation Method Specifications

	Fast Storage	**Log Monitoring**	**PageRank Computation**
Operations	put, get, delete union	put, get, filter, aggregation	get, transform, filter, order by
Patterns	single-operation	single- and multi-operation	all patterns
Data Set	randomly generated structured data	real server logs	randomly generated directed graph
Metrics	throughput	request latency statistics	test duration

2.3 Prescriptions and Prescribed Tests

Each benchmarking test is specified by a prescription. Thanks to the abstraction of processing operations and patterns, a prescribed test can be implemented over different systems for comparison. A prescription includes a subset of operations and processing patterns, a data set, a workload generation method, and the measured metrics. The subset of operation and processing patterns is selected from BigOP's whole abstraction set. The data set can be taken from real applications or generated through widely-obtainable tools. The size of the chosen data set is required to be larger than the total memory size of the SUT, so that the volume characteristic of big data is covered. The workload generation method describes how operations are issued from clients. The velocity characteristic of big data can be represented in the workload generation specification. The measured metrics can be the duration of the test, request latency statistics, and the throughput.

We instantiate three prescribed test examples in Table 3, from the simplest to the most complex. The first and the third examples[2] are taken from BigDataBench, while the second example is constructed from a common application scenario. The workload generation methods are not included in the prescription for space consideration. Instead, we describe them respectively in the following, together with an introduction to the application corresponding to each example.

Example 1. **Fast Storage**. Applications make frequent requests of data storage. This is the most basic scenario of big data acquisition. The data set is generated using BDGS' data generation tool [24]. The workload generation combines YCSB's workloads. The throughput of operations is the key metric.

Example 2. **Log Monitoring**. An application monitors its services by logs. Applications like user and server activity monitoring can be represented by this prescription. The workload generation is specified as follows. *put* is applied to some random log entries at a speed of 5000 operations per second (ops) till the total data size exceeds twice of the total system memory. Simultaneously, *get+filter+aggregation* is applied to the data set based on some random filtering condition for a *sum* result continuously.

Example 3. **PageRank Computation**. This is a core computation in the widely deployed Internet service. It is also a representative computation of graph applications. In workload generation, *get+filter+ transform* is applied to the data set *iteratively* till a given *condition* is met. The *order by* operation, a.k.a. *sort*, is executed to get the final result.

[2] Referred as Cloud OLTP and PageRank workloads respectively in BigDataBench.

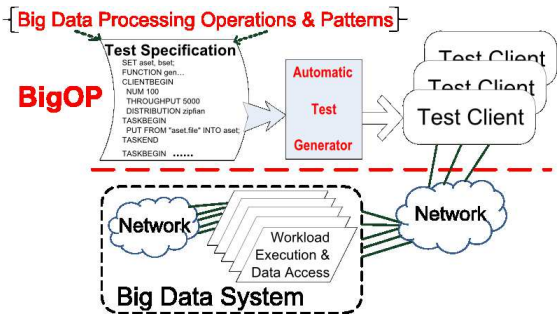

Fig. 2. Testing process and functional components of BigOP.

In constructing test prescriptions, we can adjust factors in the prescription according to our needs. For example, to test for data scalability, we can increase the workload of *put*. Similarly, to test more complicated computations, we can increase the number of combinations for *multi-operation* and *iterative* processing patterns, as well as defining a sophisticated *transform* function.

3 Evaluation

Based on the BigOP framework, we implement an automatic test generator, which can generate tests from prescriptions. Figure 2 demonstrates the testing procedure of the evaluation.

We benchmark three widely used big data processing systems, i.e. MySQL cluster (SUT1), Hadoop+HBase (SUT2) and HDFS+Spark (SUT3). All the three SUTs are deployed over five physical nodes connected with 1Gbps network. Each node has 32 GB memory and a processor with six 2.40GHz cores. The operating system is Centos release 5.5 with Linux kernel 2.6.34. Among our three examples, we only choose the second and the third for the evaluation, since the first example has been extensively tested in other benchmarks.

Log Monitoring. Figure 3 demonstrates the resulting performances of SUT1 and SUT2. SUT2 excels under the frequent record insert task as expected. SUT1 performs much better in statistics computation tasks because of the long starting time of jobs in SUT2. For complex multi-operation tasks, SUT1 is very likely to excel the others as well. The reason is as follows. Even though MapReduce-like systems including Hadoop, Shark and Hive support complex user defined

Log Monitoring—Record Insert Task

Systems	Throughput (ops/sec)	Min Latency (us)	Max Latency (us)	Average Latency (us)
MySQL Cluster	3367.795612	0	46400516	295.3059112
Hadoop+HBase	16918.51187	3	16392237	56.32552309

Log Monitoring—Statistics Task

Systems	# of Runs	Max run time(sec.)	Min run time(sec.)
MySQL Cluster	2675	5204.483333	0.008
Hadoop+HBase	6	7312	614

Fig. 3. System performances under log monitoring workloads. (The zero min latency of SUT1 is due to the batch commit mode.)

functions, they have not considered optimizations on multi-operation and iterative processing patterns. On the contrary, traditional databases are strictly SQL compliant and heavily optimized for relational queries, which usually contain single-operation and multi-operation processing patterns.

Pagerank Computation. We generate 0.5 million pages and 3.7 million links using an open-source tool [24], resulting in more than 250GB data (including page contents), which is larger than the total system memory. The running time of the task over SUT2 is ***363 seconds***, while that over SUT3 is ***96 seconds***. We also run this test over SUT1 through a stored procedure, which contains costly large-scale joins leading to intolerable test durations. The indication here is *twofold*. First, distributed in-memory computation is effective for the *iterative* computation pattern, while frequent disk accesses and network communications can be costly. Second, there is still much optimization space for distributed computation in relational database, especially when the iterative pattern is involved.

We report a small fraction of evaluation results here due to the page limit. Further benchmarking results can be found on the BigDataBench webpage[3].

Discussion. Existent big data benchmarks only cover part of BigOP's abstraction of processing operations and patterns. We take YCSB [18], TPC-DS [26], and BigBench [22] for example. YCSB is mainly for NoSQL database benchmarking. Its workload consists of only *put* and *get* operations, though which can be combined by the *multi-operation* pattern to form *scan* and *read-modify-write* operations. TPC-DS covers all abstracted operations and the first two patterns, except for the *iterative* pattern. It targets a single application domain, i.e., decision support applications. The involved data is structured data. Thus, it is not as flexible in suiting different benchmarking requirements as BigOP. BigBench extends TPC-DS. It adds new data types of semi-structured data and unstructured data, but it still does not include the *iterative* pattern.

4 Related Work

BigBench [22] is the recent effort towards designing a general big data benchmark. BigBench focuses on big data analytics, thus adopting TPC-DS as the basis and adding atop new data types like semi-/un-structured data, as well as non-relational workloads like sentiment queries. Although BigBench has a complete coverage of data types, it targets only a specific big data application scenario, not covering the variety of big data processing workloads.

The AMPLab of UC Berkeley also proposes a big data benchmark [2] in recent years. It is the systems of Spark [28] and Shark [20] that inspire the design of the benchmark, which thus targets real-time analytic applications. The benchmark not only has a limited coverage of workloads, but also covers only relational data.

Industrial players also try to develop their benchmark suites. Yahoo! release their cloud benchmark specially for data storage systems, i.e. YCSB[18]. Having its root in cloud computing, YCSB is mainly for scenarios like that in the *Fast Storage* example. The characteristics of and the diverse application workloads

[3] http://prof.ict.ac.cn/BigDataBench/

for big data are not considered in YCSB. CALDA[25] effort represents a micro benchmark for Big Data analytics. It has compared Hadoop-based analytics to a row-based RDBMS one and a column-based RDBMS one.

There exist also benchmarks targeting a specific platform, e.g. Hadoop [5]. HiBench [23] is a widely-known benchmark for Hadoop MapReduce. Hence, its four categories of workloads are limited to MapReduce based processing. It exploits stochastic methods to generate data for its workloads. However, its randomly generated data misses various features of real big data. Besides, its choice of workloads lacks of coverage, as well as solidly founded grounds. Gridmix [4] and Pigmix [10] are also two benchmarks specially designed for Hadoop MapReduce. They include a mix of workloads, including sort, grep and wordcount. They are also suffering from incomplete coverage of data and workloads.

Across different research fields, there are multiple famous and well established benchmarks. The TeraSort or GraySort Benchmark [11] considers the performance and the cost involved in sorting a large number of 100-byte records. The workload of this benchmark is too simple to cover the various needs of big data processing. SPEC [12] works well over standalone servers with a homogeneous architecture, but is not suitable for the emerging big data platforms with large-scale distributed and heterogeneous components. TPC [13] series of benchmarks are widely accepted for database testing, but only consider structured data. PARSEC [15] is a well-known benchmark for shared-memory computers, thus not suited for big data applications that mainly take a shared-nothing architecture.

5 Conclusion

BigOP is targeted at big data systems that support part of or all of its processing operation and pattern abstractions. Users of BigOP can flexibly construct tests through prescriptions for their application-specific or general benchmarking needs. Benchmarking tests can be constructed based on BigOP's abstracted operations and patterns. Big data systems that can implement the same prescribed test are comparable. System users can prescribe tests targeting at their specific application scenarios, while system developers can carry out general tests with all abstracted operations and patterns mixed and randomly generated in the workload. We design BigOP as a benchmarking framework in considering the variety of big data and its applications, which we believe is a good trade-off between benchmarking flexibility and conformity to real big data processing requirements.

Acknowledgments This work is supported in part by the State Key Development Program for Basic Research of China (Grant No. 2014CB340402) and the National Natural Science Foundation of China (Grant No. 61303054).

References

1. Amazon redshift service. http://aws.amazon.com/cn/redshift/.
2. Amplab big data benchmark. https://amplab.cs.berkeley.edu/benchmark/.
3. Cloudera impala. http://www.cloudera.com/content/cloudera/en/products-and-services/cdh/impala.html.

4. Gridmix. http://hadoop.apache.org/mapreduce/docs/current/gridmix.html.
5. Hadoop project. http://hadoop.apache.org/.
6. Hbase project. http://hbase.apache.org/.
7. Hive project. http://hive.apache.org/.
8. Hpl benchmark home page. http://www.netlib.org/benchmark/hpl/.
9. Nosql databases. http://nosql-database.org/.
10. Pigmix. https://cwiki.apache.org/confluence/display/PIG/PigMix.
11. Sort benchmark home page. http://sortbenchmark.org/.
12. Standard performance evaluation corporation (spec). http://www.spec.org/.
13. Transaction processing performance council (tpc). http://www.tpc.org/.
14. K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, et al. The landscape of parallel computing research: A view from berkeley. Technical report, UCB/EECS-2006-183, EECS Department, University of California, Berkeley, 2006.
15. C. Bienia, S. Kumar, J. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implications. In *Proc. of PACT 2008*, pages 72–81. ACM, 2008.
16. Y. Chen, F. Raab, and R. H. Katz. From tpc-c to big data benchmarks: A functional workload model. Technical Report UCB/EECS-2012-174, EECS Department, University of California, Berkeley, Jul 2012.
17. E. F. Codd. A relational model of data for large shared data banks. In *Pioneers and Their Contributions to Software Engineering*, pages 61–98. Springer, 2001.
18. B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with ycsb. In *Proc. of SoCC 2010*.
19. J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
20. C. Engle, A. Lupher, R. Xin, M. Zaharia, M. J. Franklin, S. Shenker, and I. Stoica. Shark: fast data analysis using coarse-grained distributed memory. In *Proc. of SIGMOD 2012*, pages 689–692. ACM, 2012.
21. M. Ferdman, A. Adileh, O. Kocberber, S. Volos, M. Alisafaee, D. Jevdjic, C. Kaynak, A. Popescu, A. Ailamaki, and B. Falsafi. Clearing the clouds: A study of emerging workloads on modern hardware. *Architectural Support for Programming Languages and Operating Systems*, 2012.
22. A. Ghazal, M. Hu, T. Rabl, F. Raab, M. Poess, A. Crolotte, and H.-A. Jacobsen. Bigbench: Towards an industry standard benchmark for big data analytics. In *Proc. of SIGMOD 2013*. ACM, 2013.
23. S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang. The hibench benchmark suite: Characterization of the mapreduce-based data analysis. In *Proc. of ICDEW 2010*, pages 41–51. IEEE, 2010.
24. Z. Ming, C. Luo, W. Gao, R. Han, Q. Yang, L. Wang, and J. Zhan. Bdgs: A scalable big data generator suite in big data benchmarking. *arXiv:1401.5465*, 2014.
25. A. Pavlo, E. Paulson, A. Rasin, D. J. Abadi, D. J. DeWitt, S. Madden, and M. Stonebraker. A comparison of approaches to large-scale data analysis. In *Proc. of SIGMOD 2009*, pages 165–178, New York, NY, USA, 2009. ACM.
26. M. Poess, R. O. Nambiar, and D. Walrath. Why you should run tpc-ds: a workload analysis. In *Proc. of VLDB 2007*, pages 1138–1149. VLDB Endowment, 2007.
27. L. Wang, J. Zhan, C. Luo, Y. Zhu, Q. Yang, Y. He, W. Gao, Z. Jia, Y. Shi, S. Zhang, C. Zhen, G. Lu, K. Zhan, and B. Qiu. Bigdatabench: a big data benchmark suite from internet services. Accepted by HPCA 2014.

28. M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proc. of NSDI 2012*, pages 2–2. USENIX Association, 2012.